

IN THE CLAIMS

Please cancel claims 16-26.

No amendments have been made to the remaining pending claims, which follow:

1. (Original) A method for sharing resources across a plurality of computing platforms, comprising:
 - receiving a resource access request to access a shared resource at a first computing platform;
 - determining a second computing platform via which the shared resource may be accessed;
 - sending the resource access request to the second computing platform; and
 - accessing the shared resource via the second computing platform.
2. (Original) The method of claim 1, wherein the plurality of computing platforms comprise a plurality of server blades operating in a blade server environment.
3. (Original) The method of claim 1, wherein the method is performed in a manner that is transparent to operating systems running on the plurality of computing platforms.
4. (Original) The method of claim 1, wherein the method is facilitated by firmware running on each of the plurality of computing platforms.
5. (Original) The method of claim 1, wherein the resource access request is sent to the second computing platform via an out-of-band (OOB) communication channel.

6. (Original) The method of claim 5, wherein the OOB communication channel comprises one of a system management bus, an Ethernet-based network, or a serial communication link.
7. (Original) The method of claim 5, wherein the target resource comprises a storage device.
8. (Original) The method of claim 7, wherein the resource access request comprises a storage device write request, and the method further comprises sending data corresponding to the storage device write request via the OOB communication channel.
9. (Original) The method of claim 7, wherein the resource access request comprises a storage device read request, and the method further comprises:
retrieving data corresponding to the read request from the shared resource;
and
sending the data that are retrieved back to the first computing platform via the OOB communication channel.
10. (Original) The method of claim 1, further comprising:
maintaining global resource mapping data identifying which resources are accessible via which computing platforms; and
employing the global resource mapping data to determine which computing platform to use to access the shared resource.

11. (Original) The method of claim 10, wherein a local copy of the global resource mapping data is maintained on each of the plurality of computing platforms.

12. (Original) The method of claim 10, wherein the global resource mapping data is maintained by a central global resource manager.

13. (Original) A method for sharing a plurality of storage devices across a plurality of computing platforms, comprising:

configuring the plurality of storage devices as a virtual storage volume;

maintaining a global resource map that maps input/output (I/O) blocks defined for the virtual storage volume to corresponding storage devices that actually host the I/O blocks;

receiving a data access request identifying an I/O block from which data are to be accessed via the virtual storage volume;

identifying a computing platform via which a target storage device that actually hosts the I/O block may be accessed through use of the global resource map;

routing the data access request to the computing platform that is identified;
and

accessing the I/O block on the target storage device via the computing platform that is identified.

14. (Original) The method of claim 13, further comprising:

configuring the plurality of storage devices as at least one RAID (redundant array of independent disks) storage volume;

maintaining RAID configuration mapping information that maps input/output (I/O) blocks defined for said at least one RAID virtual storage volume to corresponding storage devices that actually host the I/O blocks; and

employing the RAID configuration mapping information to access appropriate storage devices in response to read and write access requests.

15. (Original) The method of claim 14, wherein the RAID virtual storage volume is configured in accordance with the RAID-1 standard.

Claims 16-26 (Canceled).

27. (Original) A blade server system, comprising:

a chassis, including a plurality of slots in which respective server blades may be inserted;

an interface plane having a plurality of connectors for mating with respective connectors on inserted server blades and providing communication paths between the plurality of connectors to facilitate in out of band (OOB) communication channel; and

a plurality of server blades, each including a processor and firmware executable thereon to perform operations including:

receive a resource access request from an operating system running on a requesting server blade to access a shared resource hosted by at least one of the plurality of server blades;

determining a target resource host from among the plurality of server blades that hosts a target resource that may service the resource access request;

sending the resource access request to the target resource host; and

accessing the target resource via the target resource host to service the resource access request.

28. (Original) The blade server system of claim 27, wherein the operations are performed in a manner that is transparent to operating systems that may be run on the plurality of server blades.

29. (Original) The blade server system of claim 27, wherein communications between the plurality of server blades is facilitated by an out-of-band OOB communication channel.

30. (Original) The blade server system of claim 29, wherein each processor supports a hidden execution mode that is employed for facilitating communication via the OOB channel.